

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (cancelled)

2. (currently amended) The method according to claim 34 ~~claim 1~~, and comprising configuring one or more interfaces to receive input/output (IO) requests from host processors directed to specified LAs and to direct all the IO requests to the first cache and the second caches which have been configured to perform at least one of the operations of retrieving data from and storing data at the specified LAs.

3. (currently amended) The method according to claim 2, wherein the one or more interfaces comprise a mapping between the first cache and the second ~~and the one or more third~~ caches and the specified ~~first and second ranges of the~~ LAs, wherein the one or more interfaces are adapted to convert the IO requests to one or more requests and to direct the one or more requests to at least one of the first cache and the second ~~and the one or more third~~ caches in response to the mapping, and wherein detecting ~~the inability~~ the incorrect operation comprises generating a reconfigured mapping between the ~~first and the one or more third~~ second caches and the ~~first and second ranges of the~~ specified LAs, and directing the one or more requests to at least one of the ~~first and the one or more third~~ second caches in response to the reconfigured mapping.

4. - 9. (cancelled)

10. (currently amended) The method according to ~~claim 9~~ claim 39, wherein implementing the minimum redistribution comprises ~~redistributing the first and the second ranges~~ distributing the second copies using a consistent hashing function.

11. (currently amended) The method according to ~~claim 9~~ claim 39, wherein implementing the minimum redistribution comprises ~~redistributing the first and the second ranges~~ distributing the second copies using a random number function.

12. - 33. (cancelled)

34. (new) A method for managing a data storage system having caches and storage devices, comprising:

configuring a first cache, and a first storage device that is connected only to the first cache by a first coupling, to contain dirty cache-data portions and clean cache-data portions corresponding to respective device-data portions stored on the first storage device, so that a given cache-data portion and a respective given device-data portion are addressable at a given logical address (LA);

connecting second caches to respective second storage devices by respective second couplings, so that each second storage device is connected to only one second cache;

duplicating respective first copies of the dirty

cache-data portions in the second caches, and distributing first copies of the device-data portions over the second storage devices, so that each dirty cache-data portion is redundantly stored on the first cache and one of the second caches, and so that any given device-data portion is redundantly stored on the first storage device and one of the second storage devices, and so that a given-first-copy-of-the-given-dirty-cache-data-portion and a given-first-copy-of-the-given-device-data-portion are addressable at the given LA;

detecting incorrect operation of the first cache so that redundancy of the dirty cache-data portions and the device-data portions is lost; and

in response to detecting the incorrect operation, duplicating second copies of the dirty cache-data portions in the second caches and distributing second copies of the device-data portions over the second storage devices, so that the given-first-copy-of-the-given-dirty-cache-data-portion and a given-second-copy-of-the-given-dirty-cache-data-portion are on different second caches, and so that the given-first-copy-of-the-given-device-data-portion and a given-second-copy-of-the-given-device-data-portion are on different second storage devices and are addressable at the given LA, so that the redundancy of the dirty cache-data portions and the device-data portions is recovered.

35. (new) The method according to claim 34, wherein duplicating the second copies comprises duplicating the second copies while performing at least one of the operations of retrieving data from and storing data at an LA different from the given LA.

36. (new) The method according to claim 34, wherein configuring the first cache comprises connecting only the first cache to the first storage device, so that the first coupling comprises a first one-to-one coupling, and wherein connecting the second caches comprises connecting only single second caches to the respective second storage devices, so that the respective second couplings are respective second one-to-one couplings.

37. (new) The method according to claim 34, wherein at least some of the first cache and the second caches have unequal capacities, wherein duplicating the respective first copies of the dirty cache-data portions and distributing the first copies of the device-data portions comprises duplicating the respective first copies of the dirty cache-data portions and distributing the first copies of the device-data portions to achieve well-balanced loading for the data system, and wherein duplicating the second copies of the dirty cache-data portions and distributing the second copies of the device-data portions comprises duplicating the second copies of the dirty cache-data portions and distributing the second copies of the device-data portions to maintain the well-balanced loading.

38. (new) The method according to claim 34, wherein distributing the second copies comprises copying the respective first copies.

39. (new) The method according to claim 34, wherein duplicating the second copies of the dirty cache-data portions comprises implementing a minimum redistribution

50307A1

of the second copies among the second caches.

40. (new) The method according to claim 34, and comprising incorporating two or more manager processing units into at least two of the first storage device, the second storage devices, the first cache, the second caches, and operating the two or more manager processing units as a system manager in a cooperative manner.